

Impact of Transient CSMA/CA Access Delays on Active Bandwidth Measurements

Marc Portoles-Comeras
Centre Tecnologic de
Telecomunicacions de
Catalunya
Castelldefels (Barcelona),
Spain
marc.portoles@cttc.cat

Albert Cabellos-Aparicio
Universitat Politecnica de
Catalunya
Departament d'Arquitectura
de Computadors
Barcelona, Spain
acabello@ac.upc.edu

Albert Banchs
Universidad Carlos III de
Madrid
Leganes, Spain
banchs@it.uc3m.es

Josep Mangues-Bafalluy
Centre Tecnologic de
Telecomunicacions de
Catalunya
Castelldefels (Barcelona),
Spain
josep.mangues@cttc.cat

Jordi Domingo-Pascual
Universitat Politecnica de
Catalunya
Departament d'Arquitectura
de Computadors
Barcelona, Spain
jordid@ac.upc.edu

ABSTRACT

WLAN devices based on CSMA/CA access schemes have become a fundamental component of network deployments. In such wireless scenarios, traditional networking applications, tools, and protocols, with their built-in measurement techniques, are usually run unchanged. However, their actual interaction with the dynamics of underlying wireless systems is not yet fully understood. A relevant example of such built-in techniques is bandwidth measurement. When considering WLAN environments, various preliminary studies have shown that the application of results obtained in wired setups is not straightforward. Indeed, the contention for medium sharing among multiple users inherent to CSMA/CA access schemes has remarkable consequences on the behavior of and results obtained by bandwidth measurement techniques. In this paper, we focus on evaluating the effect of CSMA/CA-based contention on active bandwidth measurement techniques. As a result, it presents the rate response curve in steady state of a system with both FIFO and CSMA/CA-based contending cross-traffic. We also find out that the distribution of access delay shows a transient regime before reaching a stationary state. The duration of such transient regime is characterized and bounded. We also show how dispersion-based measurements that use a short number of probing packets are biased measurements of the *achievable throughput*, the origin of this bias lying on the transient detected in the access delay of probing packets. Overall, the results presented in this paper have several consequences that are expected to influence the design of

bandwidth measurement tools as well as to better understand the results obtained with them in CSMA/CA links.

Categories and Subject Descriptors

C.4 [Performance of Systems]: Measurement Techniques

General Terms

Theory, Measurement, Performance

Keywords

Bandwidth measurements, Wireless, CSMA/CA, Achievable Throughput

1. INTRODUCTION

WLAN devices have become a fundamental component of network deployments. They can be found in scenarios that range from simple single-hop home networks to complex mesh-like multi-radio multi-hop infrastructures. In such wireless scenarios, traditional networking applications, tools, and protocols, with their built-in measurement techniques, are usually run unchanged over wireless links. However, their actual interaction with the dynamics of underlying wireless systems is not yet fully understood.

A relevant example of such built-in techniques is bandwidth measurement. Its interest is exemplified by the many applications found in the literature, including congestion control algorithms [26, 27, 28, 31], overlay routing [29], dynamic server selection [30], and inter-domain path monitoring [22], among others. As a result, bandwidth measurements have become a mature research topic with well-developed results both at a practical level (e.g. [1, 17, 18, 19, 20, 22, 23]) and, lately, at a more fundamental level [14, 15]. However, most results have been obtained in wired environments.

When considering WLAN environments, various preliminary studies have shown that the application of these results

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IMC'09, November 4–6, 2009, Chicago, Illinois, USA.

Copyright 2009 ACM 978-1-60558-770-7/09/11 ...\$10.00.

is not straightforward ([2, 3, 25]). The main reasons for this reside in the assumptions taken to develop bandwidth measurement models and tools. In fact, traditional active bandwidth measurement techniques assume a single bit-carrier multiplexing several users in FIFO order (e.g. [1]). But, when applied to WLAN environments, this fundamental assumption does not hold any longer. In this case, the contention, among multiple users, for medium access inherent to CSMA/CA access schemes has relevant consequences on the behavior of bandwidth measurement techniques [3].

In this paper, we focus on evaluating the effect of CSMA/CA-based contention on active bandwidth measurement techniques. As a consequence, the results and conclusions derived not only apply to wireless environments, but also to any CSMA/CA-based system (e.g. PLC). Other effects appearing as a consequence of wireless channel impairments are not dealt with in this paper.

Furthermore, this paper uses an analytical framework that better accounts for the particularities of CSMA/CA links. The results of applying this framework are validated through extensive experimentation and simulation.

In short, the contributions of the paper follow:

- It presents the rate response curve in steady state of a system with two types of cross-traffic, one sharing the queue with probing traffic and the other contending for access with it.
- It reveals how the distribution process describing the access delay (i.e. the time it takes to transmit probing packets in a CSMA/CA system) is not the same for the whole probing sequence. Instead, the distribution shows a transient regime before reaching steady-state. The duration of such transient regime is characterized and bounded.
- It shows how using dispersion-based measurements to infer steady-state bandwidth metrics may suffer of deviations, specially when the number of probing packets is short. The origin of the deviations lies in the transient regime detected.

The results presented in this paper have several consequences that transcend its scope. First, we extend previous studies [25], showing that tools designed to measure available bandwidth in wired environments in fact measure *achievable throughput* in CSMA/CA links. Second, we show how the packet pair technique [23], widely used in the wireless mesh routing literature [21], constitutes a biased measure of the *achievable throughput*. Third, we introduce a simple yet effective method to improve the accuracy and convergence properties of bandwidth measurement tools based on a previous characterization of the transient regime. Interestingly, this method not only improves measurements in wireless scenarios but also in wired ones.

Furthermore, we follow a packet-based (i.e. network-layer) approach in which no knowledge of lower layer details are assumed. This approach is taken to not limit the applicability of our findings to restricted paths. Overall, these contributions are expected to help building tools that better take into account the characteristics of CSMA/CA links.

Even though they can be expected to happen, the existence and impact of transient-states when probing a system with trains of packets has not been extensively considered. An exception to this is [15] where the authors characterize,

analytically, the transients present in a FIFO queue with constant service rate and bursty cross-traffic. Following the framework developed in [15], we extend it to consider also transients present in CSMA/CA systems. Additionally, our findings are validated by means of simulation and experiments over a WLAN testbed.

The rest of the paper is organized as follows. Section 2 presents the current state-of-the-art of bandwidth measurement in the presence of CSMA/CA links. Section 3 introduces the approach to CSMA/CA links used along the paper and provides a complete steady-state rate response curve describing CSMA/CA links. Section 4 studies the transient regime of the access delay of active probing packets traversing a CSMA/CA link. Sections 5 and 6 introduce the analytical framework used to study the impact of the transitory regime and use it to derive its consequences on bandwidth measurements over CSMA/CA links. Section 7 discusses the consequences of the findings of the paper while section 8 concludes the paper. Finally, the appendix A describes the tools that have been used to validate the results.

2. BACKGROUND

The rate response curve [13, 14] is one of the basic models used in bandwidth measurement literature to characterize network paths. It essentially describes the relation between the input rate (r_i) and output rate (r_o) that a flow observes when traversing a network path. Multiple bandwidth measurement tools, specially those related to measuring the *available bandwidth*, are based on the rate response model of a FIFO queue. Such model places fluid assumptions on the cross-traffic that traverses the same FIFO queue as the probing flow and states the following relation,

$$r_o = \min(r_i, C \frac{r_i}{r_i + C - A}) = \begin{cases} r_i & r_i \leq A \\ C \frac{r_i}{r_i + C - A} & r_i \geq A \end{cases} \quad (1)$$

where C is the *capacity* or rate at which data is sent and A is the *available bandwidth* i.e. the part of C that is not being used. Further, in [14], the authors show how the rate-response curve of a FIFO queue is an accurate description of the expected interaction between the probing traffic and the cross-traffic when the system is in steady-state. However, they show how the first packets of a probing sequence are not in such steady-state conditions which may lead to measurement errors.

Recent literature related to bandwidth measurement in *wireless* networks [3, 5, 25] has reported that existing tools that aim at measuring the *available bandwidth* are not accurate. Many of these measuring tools have been designed following the rate response model of the FIFO queue. However, as shown in [25], the rate response curve for an IEEE 802.11 system differs from that of a FIFO queue. The main reason behind this is the protocol used to access the medium, the DCF, that uses a CSMA/CA mechanism to regulate medium access between contending stations. Under the DCF mechanism packets from different stations are not scheduled in FIFO order, thus breaking the assumption taken in (1).

In order to formulate a rate response curve describing WLAN links, researchers have identified the need to use different bandwidth metrics to describe their behavior. In particular, the authors of [4] propose using the *achievable throughput* metric, however they provide an empirical defi-

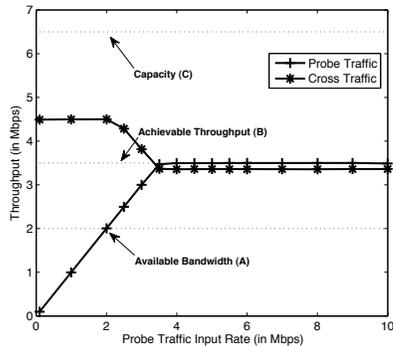


Figure 1: Experimental steady-state rate response curve of probe traffic in a WLAN setting versus throughput of cross-traffic flow. $C=6.5\text{Mbps}$, $A=2\text{Mbps}$, $B=3.4\text{Mbps}$ (tesbed)

dition to be used in IEEE 802.11 based links. Further, the authors of [25] propose the use of the concept of *fair-share* related to the functionality of the system in backlog.

Here we inherit the same term from [4] but propose, alternatively, the following definition of the *achievable throughput* B ,

$$B = \sup\{r_i : \frac{r_o}{r_i} = 1\} \quad (2)$$

The reason behind using this definition will be made clear later but notice, in advance, that in (1) the *achievable throughput* B corresponds to the *available bandwidth* A .

Now, with the *achievable throughput* metric, the authors of [25] propose the following rate response curve to describe the behavior of a probing flow that contends for channel access in an IEEE 802.11 system.

$$r_o = \min(r_i, B) \quad (3)$$

In this case the *achievable throughput* corresponds to the *fair-share* portion of the capacity that the probing flow can get when contending for channel access with other wireless stations. Note that, according to (3), the *available bandwidth* can only be detected when it coincides with B , and this only happens under certain conditions in CSMA/CA systems.

In order to illustrate all this consider the experimentation result depicted in figure 1. The figure plots the rate response curve describing the interaction of a probing flow contending for channel access with another flow¹. The figure also shows the evolution of the cross-traffic throughput for each probing rate. As it can be seen, when the cross-traffic starts experiencing a decrease in its throughput, that is, when the probing traffic arrives at the *available bandwidth* ($\sim 2\text{Mbps}$), the rate response curve shows no sign of deviation. Instead, the rate response curve flattens when the probing rate reaches the *fair-share* ($\sim 3.5\text{Mbps}$) that it can get from the wireless medium. This *fair-share* corresponds, in fact, to the *achievable throughput* defined above.

The present paper completes this analysis taking two basic observations of the system. First, we notice that cross-traffic may not only appear in the access but can also share

¹In order to obtain the rate response curve we use long probing trains (>10000 packets) and evaluate it in steady-state

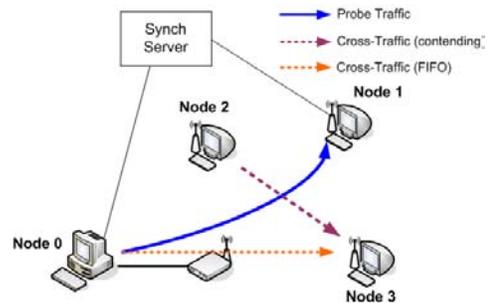


Figure 2: Experimental/simulation scenario

transmission queue with probing traffic. Second, we show that in a CSMA/CA access the interaction between probing traffic and the system presents a transient in the delay to access the medium. This transient is not present in (*wired*) FIFO systems and produces different deviations than those described in [14].

2.1 Validation Setup

The study presented in this paper is based on theoretical analysis, simulation and experimentation. In order to validate our model we have reproduced it (figure 2) in a testbed, simulator (NS2) and a Matlab queing simulator (see figure 3). The interested reader can find all the details in the appendix A. It is worth noting here that unless noted otherwise, the results presented in this work have been obtained from repeating experiments over 80 times while the simulations have been repeated 25.000 (NS2) to 70.000 (Matlab) times. Further, the cross-traffic generated follows a Poisson distribution.

3. RATE RESPONSE CURVE IN STEADY STATE: COMPLETE PICTURE

3.1 Model of a WLAN link

A considerable part of bandwidth measurement studies over wireless networks consider only the inter-relation between the probing flow and the access to the wireless medium. However, one should also consider the possibility that the station that is being used to measure is also transmitting data at the same time. In such a case probing packets would be sharing the transmission queue with such traffic before entering contention for channel access. As a consequence, the probing flow can interact with cross-traffic in two different ways when it traverses a WLAN link. Figure 3 illustrates this fact.

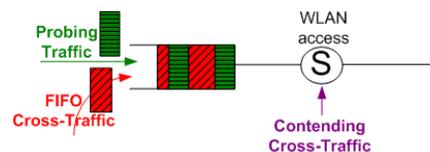


Figure 3: Model of the interaction between probing traffic and cross-traffic in a WLAN system

On one side the probing flow shares the transmission queue with cross-traffic that the wireless station sends at the same time. We refer to this type of cross-traffic as *FIFO cross-traffic* throughout the rest of the paper. On the other side, once a probing packet is at the head of the transmission queue it has to contend for channel access with the *contending cross-traffic* from other wireless stations.

The interaction between probing traffic and *contending* cross-traffic is not considered at a bit or packet per second point of view but in terms of access delay. The paper takes the perspective that there is a 'virtual scheduler' (S in the figure) that regulates channel access (in the context of the paper it follows a CSMA/CA mechanism) and delays packets according to a certain scheduling rule. As will be further developed later, we are interested in knowing the characteristics of the access delay of probing packets. That is, the delay since they are at the head of the transmission (FIFO) queue until they are completely transmitted (i.e. scheduling + transmission time).

3.2 The rate response curve in steady state

This section extends the rate response curve of a CSMA/CA system to account for both types of interaction between the probing flow and cross-traffic described above.

Equation (3) shows that the *fair-share* that the probing traffic can get out of the wireless medium constitutes a limiting bound to its transmission rate. The interaction between FIFO cross-traffic and the probing flow, reduces then, in steady-state, to a FIFO interaction as described in equation (1) but with the *fair-share* acting as the bandwidth to share between probing and FIFO cross-traffic flows. The following expression describes this interaction and constitutes the rate response curve, in steady-state, of the system considered.

$$r_o = \begin{cases} r_i & r_i \leq B \\ B_f \frac{r_i}{r_i + u_{fifo} B_f} & r_i \geq B \end{cases} \quad (4)$$

In this expression B_f represents the *achievable throughput* that the probing flow would get if there is no FIFO cross-traffic and u_{fifo} is the mean portion of time that the FIFO cross-traffic is using the system. Further, the *achievable throughput* B can be expressed here as,

$$B = B_f(1 - u_{fifo}) \quad (5)$$

Figure 4 is an experimental illustration of expression (4). The plot shows how the rate response curve starts deviation when the aggregate FIFO cross-traffic and probing traffic achieve the fair-share that the wireless station can get out of the wireless medium. After that, as the probing traffic increases its throughput it gains presence in the FIFO queue in detriment of the FIFO cross-traffic.

4. TRANSIENT-STATE BEHAVIOR OF THE ACCESS DELAY

This section analyzes the characteristics of the access delay process describing the time that packets wait at the head of the FIFO transmission queue until they gain channel access and are completely transmitted. For this study we remove any FIFO cross-traffic from the proposed model and focus, strictly on the interaction between probing traffic and the *contending* cross-traffic. Figure 5 illustrates the scenario considered here.

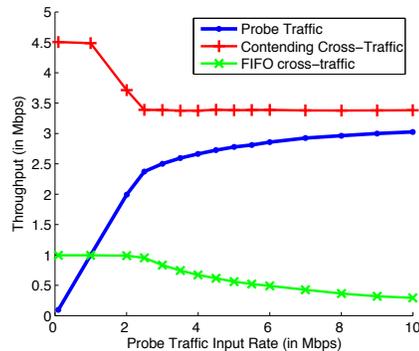


Figure 4: The complete picture

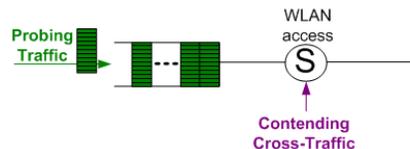


Figure 5: Model of the interaction between probing traffic and (contending) cross-traffic in a WLAN system

The access delay in CSMA/CA based systems has been repeatedly studied in the literature. Indeed, different researchers have analyzed its exact distribution using Markov Chains [6, 8]; others show how the exponential distribution provides a good fit [7]. All these studies focus on modeling the steady-state distribution of the access delay. However, in general, active bandwidth measurements are gathered using probing trains of a limited number of packets in order to limit intrusiveness. As a consequence, for the purpose of this work, we are interested in analyzing how the access delay evolves over time as an increasing number of packets are sent. In other words, we are interested in the transient-state behavior of the access delay in the system we are considering.

In order to illustrate this evolution first consider the following experiment: using NS2 we send 1000 probe packets

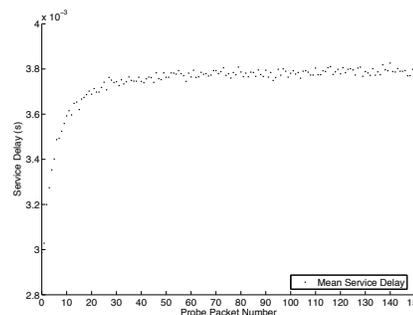


Figure 6: Mean access delay vs. Probe packet num (simulator)

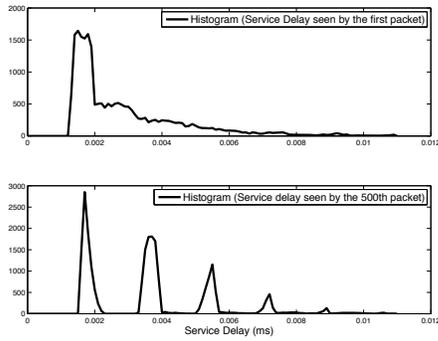


Figure 7: Histogram of the s.d seen by the first and 500th packet (simulator)



Figure 8: Analysis of the distribution (8Mbps probe-traffic rate, 2Mbps cross-traffic rate (Top) KS-Test (Bottom) Mean queue size (simulator)

at a given rate (5Mbps) and with a given load of contending cross-traffic (4Mbps). We have repeated the experiment 25000 times and, for each probe packet (indexed from 1 to 1000), we compute the distribution of the access delay individually (considering all the repetitions).

Figure 6 plots the average access delay that each one of the first 150 packets observes. The figure shows how the average access delay perceived by the first packets is lower than for the rest of them. This suggests that, in fact, the distribution of the access delay changes as more probe traffic keeps on arriving to the WLAN link. In order to verify this hypothesis, figure 7 plots the histogram of the access delay as seen by the first probe packet and by the 500th. As the plot shows, the distribution changes significantly. The main reason behind this is that as new probing packets keep on arriving they keep on increasing the load of the network until reaching a steady-state of interaction with the (contending) cross-traffic.

To further illustrate this we use the well-known Kolmogorov-Smirnov² (KS) goodness-of-fit test [16]. This statistical test

²Since we are using the KS test to compare two empirical discrete distributions we convert one of them to a continuous one using linear interpolation.

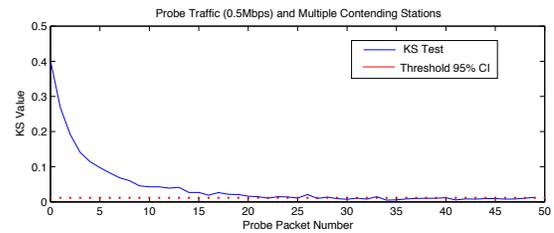


Figure 9: Analysis of the distribution (complex case) (simulator)

is used to compare the resemblance of the delay distribution suffered by every probing packet starting from the first one, and the delay distribution once probing packets have reached a steady-state. The KS test is non-parametric and analyzes whether two different sets come from the same random distribution. Using this test we compare the distribution of each individual packet in the probing sequence with the access delay distribution of the last 500 probing packets.

Figure 8 -top- shows the evolution of the KS-test of the distribution of each one of the first 100 probing packets when compared to steady-state distribution. The figure reveals clearly how there is a transient-state in the access delay that the probing packets suffer when contending for channel access. The KS-test evolution is put in contrast to the evolution of the mean queue size of a cross-traffic station that contends for channel access (see figure 8 -bottom-). The comparison shows how the transient-state duration follows the amount of time that it takes to the contending queue to reach a stationary size (~10 packets).

We have also experimented with more complex scenarios. As an example consider figure 9 that shows the KS-test for a case with 4 contending stations using different packet sizes (40, 576, 1000 and 1500 bytes) and the following rates respectively (0.1, 0.5, 0.75 and 2Mbps). Again, the figure reveals a transitory regime in the distribution of the access delay, also load of the system before and after the probing flow enters the system. As the figure shows, we need to send tens of packets until reaching a steady-state.

We have simulated more cases with different degrees of complexity obtaining similar results. The transient-state is present whenever the system is not empty, nor in backlog when the probing flow starts being transmitted. The transitory is maximum when either probing traffic and/or contending traffic are exactly sending at their *fair-share* of the wireless medium. Finally, it is worth noting here that the system does not need to reach saturation to show a transient-state behavior but it also appears in unsaturated conditions.

4.1 Duration of the transient-state of access delay

In order to design efficient measurement strategies over CSMA/CA systems we need to provide some bounds on the duration of the transient-state of the access delay. The main hypothesis taken here is that the duration of the transient-state has a close relation to the *offered traffic load* that both probing and contending stations are trying to inject into the network.

Figure 10 plots a simulation based estimation of the duration of the transitory. To generate the figure we have fixed an offered probing load of 1 Erlang. The transitory is es-

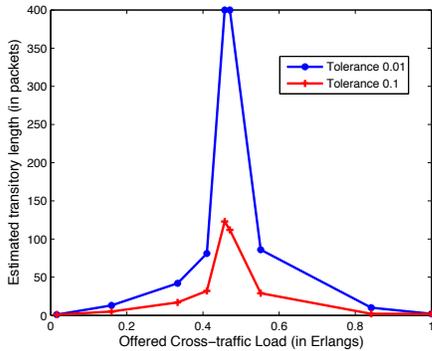


Figure 10: Estimated duration of the transitory with 0.1 and 0.01 tolerance. Offered probing load=1Erlang.

timated for a range of values of offered cross-traffic load. In order to estimate the duration of the transitory we have conducted extensive repetitions of the simulation in order to assure proper convergence of access delay distributions. The figure identifies the first packet, for each cross-traffic load, whose average access delay is within 0.1 or 0.01 of the steady-state average value.

The offered cross-traffic load at which the length of the transient-state is maximum corresponds in fact to its *fair-share*. This not only applies to the cross-traffic load but also to the probing load. When the offered load of the probing flow corresponds to its *fair-share* the duration of the transient-state is also maximum.

In order to provide practical values of the transitory we have conducted an extensive simulation for a range of probing and cross-traffic loads with multiple contending stations. In order to determine the length of the transitory we find the first packet whose average access delay lays within 0.1 of the expected access delay in steady-state conditions. We have found that, under this conditions, the transient-state never exceeds 150 packets.

4.2 Consequences of the observations

This section has shown and characterized the transient-state behavior of the access delay of probing packets when traversing a CSMA/CA link. The results in this section imply that the first packets of a probing sequence do not capture the long-term behavior of larger flows but represent deviated samples of the steady-state interaction between the probing flow and cross-traffic. This observation has a direct impact on bandwidth measurement tools that generally use short trains of packets to support measurements.

5. MODELING PACKET DISPERSION WITH TRANSIENT ACCESS DELAY

5.1 Analytical framework

Here we introduce the basic analytical framework used to deal with this problem. This framework was originally proposed in [14] but is extended here to include the particularities of WLAN transmissions. We have chosen to use this framework as it has been designed to capture the effects of possible transient-state on dispersion measurements. Fur-

thermore, the analytical framework has been adapted to the particularities of WLAN transmissions and extended to capture the transitory evolution of the access delay.

5.1.1 Sample-path analysis

The reference framework used here (and proposed in [14]) uses sample-path arguments to analyze rate response curves of dispersion measurements. This avoids the probabilistic treatment of the processes involved and generalizes results obtained to most practical cases.

This paper follows the same approach and adopts the sample-path arguments from the original framework. We believe that this facilitates reading through the arguments exposed throughout sections 5 and 6.

We also inherit from the reference framework the use of the expression $E[\cdot]$ to denote the limiting average of a sample-path of a process. Thus in the rest of the paper we use the following two definitions for continuous and discrete-time sample-paths,

$$E[Q(t)] = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t Q(u) du$$

$$E[Q] = \lim_{m \rightarrow \infty} \frac{1}{m} \sum_{k=1}^m Q^{(k)}$$

In the second expression, for discrete-time sample-paths, m refers to the number of measurement sequences in a sample-path, and k refers to the k th sequence.

5.1.2 The probing sequence: Arrivals, departures and input gap

The probing sequence consists of a series of n packets that enter the transmission queue at instants $\{a_i, i = 1, 2, \dots, n\}$. Their departure instants, meaning the time at which they are completely transmitted, form the sequence $\{d_i, i = 1, 2, \dots, n\}$. We are considering here periodic probing flows with a fixed inter-packet arrival time or input gap: $g_I = a_i - a_{i-1}$. Finally, in order to obtain the measure a number m of probing sequences are sent through the target link. These probing sequences are sent with Poisson spacing between them in order to assure complete interaction with the system.

5.1.3 The access delay process

The access delay that probing packets experience is a random process. This process is the result of the interaction between probing traffic, contending cross-traffic and back-off. To account for this let us define the sequence $\{\mu_i, i = 1, 2, \dots, n\}$ to denote the access delay that each one of the n probing packets of a probing sequence experiences when contending for medium access.

5.1.4 Processes associated to cross-traffic in the FIFO part of model

The cross-traffic that shares the transmission queue with probing traffic comes associated to the *hop-workload process* $\{W(t), 0 \leq t < \infty\}$, that is the sum of the service times of all cross-traffic packets in the FIFO queue and the remaining service time of any cross-traffic packet that may be in service. Note that cross-traffic packets experience also a random access delay implicit in the hop-workload process. Note also that this process refers to the cross-traffic only, without considering the probing flow.

Taking into account the *hop-workload process*, the utilization of the FIFO queue can be defined such that,

$$U(t) = \begin{cases} 1 & W(t) > 0 \\ 0 & W(t) = 0 \end{cases} \quad (6)$$

We are interested in the cumulative use of the FIFO queue along time, that we define as

$$u_{fifo}(t) = \frac{1}{t} \int_0^t U(s) ds \quad (7)$$

Further, this paper assumes workload stability during the measurement process. We use the term \bar{u}_{fifo} to denote the limiting average utilization that the cross-traffic makes of the FIFO queue,

$$\bar{u}_{fifo} = E[u_{fifo}(t)] \quad (8)$$

With some abuse of notation we use the term $u_{fifo}(t, t+\tau)$ to refer to the average utilization of the queue during the period $(t, t+\tau]$. That is,

$$u_{fifo}(t, t+\tau) = \frac{1}{\tau} \int_t^{t+\tau} u_{fifo}(u) du \quad (9)$$

Finally, we define the *offered workload* of cross-traffic $\{X(t), 0 \leq t < \infty\}$ as the cumulative sum of the service time of cross-traffic packets that enters the FIFO queue. We define also the averaging function $\{Y(t, t+\tau), 0 \leq t < \infty\}$ as,

$$Y(t, t+\tau) = \frac{X(t+\tau) - X(t)}{\tau} \quad (10)$$

Note that under the assumption of stability for the system,

$$E[Y(t, t+\tau)] = E[u_{fifo}(t)] = \bar{u}_{fifo} \quad (11)$$

5.1.5 Adding probe traffic in the queue: Intrusion residual

First, we define the utilization of the fifo queue $\tilde{u}_{fifo}(t, t+\tau)$ coming from the superposition of fifo cross-traffic and probing traffic.

Second, we define as $\tilde{W}(t), 0 \leq t < \infty$ the hop-workload process resulting from the superposition of the service time of the FIFO cross-traffic and that of the probing traffic. The intrusion residual $W_d(t)$ accounts for the sum of the service time of all probing packets in the FIFO queue and the remaining time to service any probing packet that may be in transmission. The intrusion residual is, then, the difference between $\tilde{W}(t)$ and $W(t)$ at any time,

$$W_d(t) = \tilde{W}(t) - W(t) \quad (12)$$

Next, we define the series $\{R_i, i = 1, 2, \dots, n\}$ which captures the intrusion residual that every probing packet finds when it enters the transmission queue³,

$$R_i(a_1) = W_d(a_i^-) = W_d(a_1 + (i-1)g_I^-) \quad (13)$$

³The minus superscript refers to the state of the queue just before the arrival, i.e. without counting the new arrival.

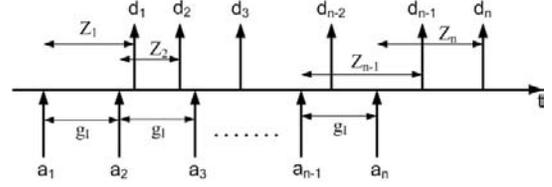


Figure 11: Inter-relationship between probing arrival sequence (a_i), departure sequence (d_i) and cross-traffic related processes (Z_i).

Note that R_i is a recursive process that under the assumptions in this work can be expressed as,

$$R_i = \begin{cases} 0 & i = 1 \\ \max(0, \mu_{i-1} + R_{i-1} - (1 - u_{FIFO}(a_{i-1}, a_i))g_I) & i > 1 \end{cases} \quad (14)$$

Finally, we define the series $\{Z_i, i = 1, 2, \dots, n\}$ that encloses the queuing plus access delay that each one of the probing packets experiences. Under the assumptions taken,

$$Z_i = d_i - a_i = \mu_i + R_i + W(a_i) \quad (15)$$

5.2 Dispersion based measurements: The output gap and its relation to the probing rate

Dispersion based measurements of bandwidth metrics consist on measuring the dispersion (or inter-departure time) of packets at the output of a path (receiving side). This measure is then used to infer the value of bandwidth related metrics. The output gap (or dispersion) of a train of probing packets is defined as follows,

$$g_O = \frac{d_n - d_1}{n - 1} \quad (16)$$

Figure 11 illustrates the contribution of the processes defined above to the value of the output gap. From the arrival of the first probing packet at the transmission queue (a_1), probing packets keep on arriving at a constant interval of g_I . The cross-traffic, access delay and the intrusion residual of previous probing packets (Z_i) randomize the departure times of probing packets (d_i) and thus, their output dispersion (g_O).

Observing figure 11 we can obtain the output gap in relation to the different processes involved.

$$g_O = \frac{d_n - d_1}{n - 1} = \frac{(n-1)g_I + Z_n - Z_1}{n - 1} \quad (17)$$

Expanding this expression we get the following,

$$g_O = g_I + \frac{R_n}{n-1} + \frac{W(a_n) - W(a_1)}{n-1} + \frac{\mu_n - \mu_1}{n-1} \quad (18)$$

The output dispersion can also be formulated taking into consideration the *offered workload* of probing traffic as,

$$g_O = \frac{1}{n-1} \sum_{i=2}^n (\mu_i + Y(a_{i-1}, a_i)) + (1 - \tilde{u}_{fifo}(d_1, d_n))g_I \quad (19)$$

The intuition behind this last expression is as follows. We take the departure of first packet as a reference (see d_1 in figure 11). The time elapsed until d_n comes from the addition of the (1) access delay of all probing packets (from second to last), (2) the workload of FIFO traffic that arrives in between probing arrivals and (3) the remaining "idle" time that is not being used for transmissions by either process.

5.3 Problem formulation

We are interested in studying whether dispersion measurements can be used to estimate the steady-state rate response curve of a wireless link with CSMA/CA access. Measurement tools based on dispersion take the assumption that the relation between the input (g_I) and output (g_O) dispersions of a probing train can be used as estimators of the interrelation between input (r_i) and output (r_o) rates of a flow traversing the system. In other words if L is the length of the packets used for probing, dispersion based measurements assume that L/g_I is a good approximation of r_i and L/g_O is a consistent estimator of r_o .

Reformulating equation (4) from a dispersion perspective, the problem of bandwidth measurement follows,

$$E[g_O] \stackrel{?}{=} \begin{cases} g_I & g_I \geq \frac{L}{B} \\ \frac{L}{B_f} + \bar{u}_{fifo} g_I & g_I \leq \frac{L}{B} \end{cases} \quad (20)$$

As (20) states we are interested in analyzing the limiting average of the output dispersion ($E[g_O]$).

First, taking the limiting average of equation (18) we obtain,

$$E[g_O] = g_I + \frac{E[R_n]}{n-1} + \kappa(n) \quad (21)$$

$$\text{with } \kappa(n) = \frac{E[W(a_n) - W(a_1)]}{n-1} + \frac{E[\mu_n] - E[\mu_1]}{n-1}.$$

Further, taking the limiting average of (19) we get a second expression for the output dispersion,

$$E[g_O] = \frac{1}{n-1} \sum_{i=2}^n (E[\mu_i] + \bar{u}_{fifo} g_I) + E[(1 - \bar{u}_{fifo}(d_1, d_n))g_I] \quad (22)$$

Expressions (21) and (22) will be used to derive upper and lower bounds to the output dispersion.

6. RATE RESPONSE CURVES IN PRESENCE OF TRANSIENT ACCESS DELAYS

This section presents an analysis of the output dispersion when probing a system with CSMA/CA access. The study provides bounds on this value rather than closed form expressions.

The basic finding here is that, when using limited number of probing packets, the transitory stage in the access delay induces deviations from the steady-state response curve. Further, these deviations are, in some sense, opposite to the ones caused by the FIFO cross-traffic itself (as detected previously in [15]). The reason behind this is that first packets are 'accelerated' in contrast to packets sent in steady-state. This leads, in some cases, to infer optimistic values of bandwidth metrics.

6.1 Part I: Analysis of the expected output dispersion

6.1.1 Intrusion residual

On one side, from expression (21), we learn that the expected output gap depends on the expected value for the residual that the last packet of the probing train (i.e. with index n) finds in the queue. Recalling the recursive expression (14) there can be defined the following bounds for the intrusion residual.

$$\max(0, \sum_{i=1}^{n-1} (\mu_i - g_I)) \leq R_n \leq \sum_{i=1}^{n-1} \mu_i \quad (23)$$

The lower bound comes from the assumption that the probing sequence did not find any cross-traffic in the FIFO queue. The upper bound considers that the probing sequence found the FIFO queue with sufficient workload so that all probing packets find each other in the queue before transmission.

Taking the average of R_n , we can differentiate two regions,

$$\frac{E[R_n]}{n-1} = \begin{cases} \frac{\beta_n}{n-1} & g_I \leq \frac{1}{n-1} \sum_{i=1}^{n-1} (E[\mu_i]) \\ \frac{\alpha_n}{n-1} & g_I \geq \frac{1}{n-1} \sum_{i=1}^{n-1} (E[\mu_i]) \end{cases} \quad (24)$$

The specific values of α_n and β_n depend on the specific characteristics of the random cross-traffic (contending and FIFO) and are bounded as follows,

$$\begin{cases} \frac{1}{n-1} \sum_{i=1}^{n-1} (E[\mu_i] - g_I) \leq \frac{\beta_n}{n-1} \leq \frac{1}{n-1} \sum_{i=1}^{n-1} (E[\mu_i]) \\ 0 \leq \frac{\alpha_n}{n-1} \leq \frac{1}{n-1} \sum_{i=1}^{n-1} (E[\mu_i]) \end{cases} \quad (25)$$

On the other side, from expression (22), we can see the dependence of the output dispersion on the time that the wireless medium is being used considering the superposition of probe traffic and FIFO cross-traffic. We can bound this value as follows,

$$\min(1, \frac{1}{g_I} \frac{1}{n-1} \sum_{i=2}^n (E[\mu_i]) \leq \bar{u}_{fifo}(d_1, d_n) \leq 1 \quad (26)$$

Note that when $g_I \leq \frac{1}{n-1} \sum_{i=2}^n (E[\mu_i])$ the FIFO queue is being used during the whole measurement process (i.e. $\bar{u}_{fifo}(d_1, d_n) = 1$), regardless of the amount of FIFO cross-traffic in the queue.

6.1.2 Bounds for the expected output dispersion

Now we reconsider expressions (21) and (22) taking into account the bounds (25) and (26) derived for the residual processes.

When $g_I \leq \frac{1}{n-1} \sum_{i=2}^n E[\mu_i]$, equations (22) and (26) provide a closed form expression for the output dispersion,

$$E[g_O] = \frac{1}{n-1} \sum_{i=2}^n (E[\mu_i] + \bar{u}_{fifo} g_I) \quad (27)$$

When $g_I \geq \frac{1}{n-1} \sum_{i=2}^n E[\mu_i]$ the output dispersion can be bounded as follows,

$$\begin{cases} \max(g_I + \kappa(n), \frac{1}{n-1} \sum_{i=2}^n (E[\mu_i] + \bar{u}_{fifo} g_I)) \leq E[g_O] \\ \min(g_I + \frac{1}{n-1} \sum_{i=1}^{n-1} E[\mu_i] + \kappa(n), (\bar{u}_{fifo} + 1)g_I) \geq E[g_O] \end{cases} \quad (28)$$

Rearranging the lower bound in expression (28) we can differentiate two regions.

$$E[g_O] \geq \begin{cases} g_I + \kappa(n) & g_I \geq \frac{\frac{1}{n-1} \sum_{i=2}^n (E[\mu_i]) - \kappa(n)}{1 - \bar{u}_{FIFO}} \\ \frac{\frac{1}{n-1} \sum_{i=2}^n (E[\mu_i]) + \bar{u}_{FIFO} g_I}{1 - \bar{u}_{FIFO}} & g_I \leq \frac{\frac{1}{n-1} \sum_{i=2}^n (E[\mu_i]) - \kappa(n)}{1 - \bar{u}_{FIFO}} \end{cases} \quad (29)$$

We can do the same for the upper bound, that presents three differentiated regions

$$E[g_O] \leq \begin{cases} g_I + \frac{1}{n-1} \sum_{i=1}^{n-1} E[\mu_i] + \kappa(n) & g_I \geq \frac{\frac{1}{n-1} \sum_{i=1}^{n-1} (E[\mu_i]) + \kappa(n)}{\bar{u}_{FIFO}} \\ (\bar{u}_{FIFO} + 1)g_I & \frac{1}{n-1} \sum_{i=2}^n E[\mu_i] \leq g_I \\ & \leq \frac{\frac{1}{n-1} \sum_{i=1}^{n-1} E[\mu_i] + \kappa(n)}{\bar{u}_{FIFO}} \\ \frac{1}{n-1} \sum_{i=2}^n (E[\mu_i] + \bar{u}_{FIFO} g_I) & g_I \leq \frac{1}{n-1} \sum_{i=2}^n (E[\mu_i]) \end{cases} \quad (30)$$

Expressions (29) and (30) give upper and lower bounds of the rate response curve of the system in transient-state. Next section provides insights into the inter-relation between them and the rate response curves in steady-state. This will help understand rate response curves obtained using probing trains with a limited number of packets.

6.2 Part II: Results without FIFO cross-traffic

This section assumes that no cross-traffic is present in the FIFO queue and analyzes the rate response curve in transient-state. The objective is to provide insights into the results obtained, for example, in [3, 5, 25].

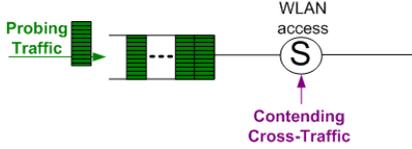


Figure 12: The system without FIFO cross-traffic

6.2.1 The achievable throughput

In this case, probing packets cannot be sent, in average, faster than $\frac{1}{n} \sum_{i=1}^n (E[\mu_i])$. As a result the achievable throughput in this particular case can be defined as,

$$\frac{L}{B} = \frac{1}{n} \sum_{i=1}^n (E[\mu_i]) \quad (31)$$

Note that this expression encloses the transient-state behavior of the access delay. Note also that as the number of probing packets grows the access delay eventually reaches a steady-state with an average access delay such that,

$$\frac{L}{B} \xrightarrow{n} E[\mu_n] \quad (32)$$

6.2.2 Bounds to the output dispersion and discussion

We rewrite here (29) and (30) for this study case. Notice that in this particular case $\kappa(n) = \frac{E[\mu_n] - E[\mu_1]}{n-1}$

$$E[g_O] \geq \begin{cases} g_I + \frac{E[\mu_n] - E[\mu_1]}{n-1} & g_I \geq \frac{1}{n-1} \sum_{i=1}^{n-1} E[\mu_i] \\ \frac{\frac{1}{n-1} \sum_{i=2}^n E[\mu_i] + \bar{u}_{FIFO} g_I}{1 - \bar{u}_{FIFO}} & g_I \leq \frac{1}{n-1} \sum_{i=1}^{n-1} E[\mu_i] \end{cases} \quad (33)$$

$$E[g_O] \leq \begin{cases} g_I & g_I \geq \frac{1}{n-1} \sum_{i=2}^n E[\mu_i] \\ \frac{1}{n-1} \sum_{i=2}^n E[\mu_i] & g_I \leq \frac{1}{n-1} \sum_{i=2}^n E[\mu_i] \end{cases} \quad (34)$$

There are a number of observations that can be done here. First considering that, as shown in section 4, the access delay μ_i is an increasing function with respect to the packet index i , the following is true for any value of $n > 2$,

$$\frac{1}{n-1} \sum_{i=1}^{n-1} (E[\mu_i]) \leq \frac{1}{n-1} \sum_{i=2}^n (E[\mu_i]) \leq E[\mu_n] \quad (35)$$

As a result, we can see that in both expressions (33) and (34) the input rate acting as a 'knee' separating different regions of the curve is higher than the (steady-state) *achievable throughput*.

However, taking into account the lower bound (33) we can observe the following. When probing at rates such that $g_I \geq \frac{1}{n-1} \sum_{i=1}^{n-1} E[\mu_i]$, the output dispersion may deviate as much as $\frac{E[\mu_n] - E[\mu_1]}{n-1}$. In other words, when the access delay variation during the transient-state is sufficiently high (in contrast to the probing rate), the output gap ($E[g_O]$) deviates (is higher) than the steady-state curve.

Finally, notice that when the probing rate is high enough (i.e. when $g_I \leq \frac{1}{n-1} \sum_{i=1}^{n-1} E[\mu_i]$) the output dispersion is higher than the one in steady-state. In other words, if we probe the system at a sufficiently high rate we might infer optimistic values of the steady-state response curve.

6.2.3 Experimental results

Figure 13 plots an experimental result illustrating these observations. The rate response curves plotted correspond to those of packet trains probing a CSMA/CA link at different rates. The figure clearly illustrates the above observations:

- First, the rate response curve follows the steady-state curve until a certain point (~ 2.5 Mbps) when the inferred rate response is lower than the steady-state one.
- Second, when probing at high rates the transient-state curves tend to overestimate the *achievable throughput* that can be achieved in steady-state.

6.3 Part III: Reintroducing FIFO cross-traffic. The complete model

6.3.1 The achievable throughput

We can define again a relation between the achievable throughput and the access delay that probing packets see (see figure 14).

$$\frac{L}{B} = \frac{\frac{1}{n} \sum_{i=1}^n (E[\mu_i])}{1 - \bar{u}_{FIFO}} \quad (36)$$

Note again that as the number of probing packets grows the limiting average of the access delay becomes constant and we can say that,

$$\frac{L}{B} \xrightarrow{n} \frac{E[\mu_n]}{1 - \bar{u}_{FIFO}} \quad (37)$$

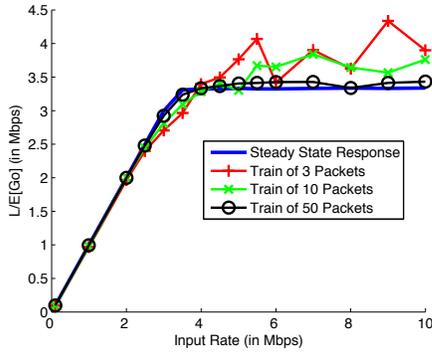


Figure 13: Experimental rate response curve of a system without FIFO cross-traffic (testbed)

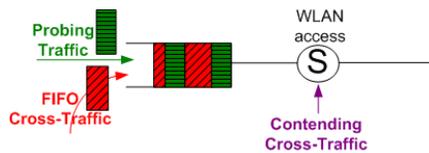


Figure 14: Model of the interaction between probing traffic and cross-traffic in a WLAN system

6.3.2 Bounds to the output dispersion and discussion

In this specific case, expressions (29) and (30) cannot be reduced. Putting in contrast these expressions with the ones taken in the previous simplified case (i.e. without FIFO cross-traffic), there are a number of observations to make.

First, the burstiness of the FIFO cross-traffic leads to looser bounds than before. As a consequence depending on the characteristics of the cross-traffic flow it is reasonable to expect higher deviations from the steady-state curve. This is specially true when probing at lower rates than the steady-state achievable throughput. As the burstiness of cross-traffic flow increases so will do the variability of dispersion measures, thus leading to higher deviations from the steady-state behavior.

Second, it can be seen that when probing the system at high rates dispersion measurements based on short packet trains tend to overestimate the steady-state rate response curve. Even more, the last region in expression (30) assures that, no matter how bursty the FIFO cross-traffic is, when probing at a sufficient rate the output dispersion will overestimate the steady-state behavior.

6.3.3 Experimental results

Figure 15 illustrates these observations. As predicted, the rate response curve inferred using packet dispersion measurements follows the steady-state behavior until it deviates (~ 2 Mbps) before reaching the *achievable throughput*. When probing at higher rates the figure shows that dispersion measurements continue overestimating the steady-state response regardless of the presence of FIFO cross-traffic.

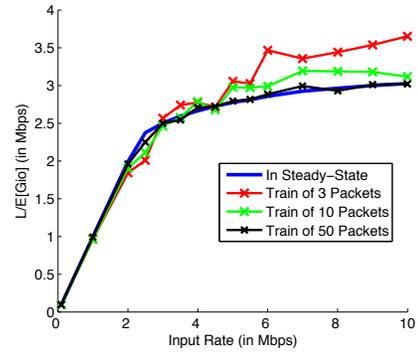


Figure 15: Experimental rate response curve for the complete system (testbed)

7. DISCUSSION ON CONSEQUENCES AND APPLICATIONS OF FINDINGS

This section discusses the main findings of this study and some consequences and possible applications that they entail.

7.1 Summary of findings

- In section 3 the paper provides a complete steady-state rate response curve of a system with CSMA/CA medium access. It takes into consideration that the probing and cross-traffic flows can interact in two differentiated manners: (1) sharing a FIFO queue and (2) contending for channel access in a random manner following the CSMA/CA procedure.
- In section 4 the study provides some insights into the transient-state behavior of CSMA/CA systems. The study shows how the duration (in number of packets) of the transient-state relates to the offered load coming from both probing and cross-traffic. We show how including some tolerance in the measurement process allows reducing the transient-state duration to values that can be used in practice when designing measurement tools.
- Finally, section 6 analyzes the impact that the transient-state evolution the access delay has on the accuracy of dispersion based measurements. It shows how dispersion measurements based on using short packet trains deviate from steady-state behavior which may lead to erroneous inferences on bandwidth metrics.

7.2 A consequence: bandwidth estimation in WLAN links

The rate response curve for FIFO queues or some of the ideas that it encloses has been repeatedly used in the literature to design bandwidth measurement tools. As defined here, the *achievable throughput* corresponds to the *available bandwidth* when applied to FIFO queues. However, as we have seen, when applied to CSMA/CA systems, the *achievable throughput* and *available bandwidth* only coincide under certain conditions.

From the results in this paper it can be argued that a large portion of the tools used to infer the *available band-*

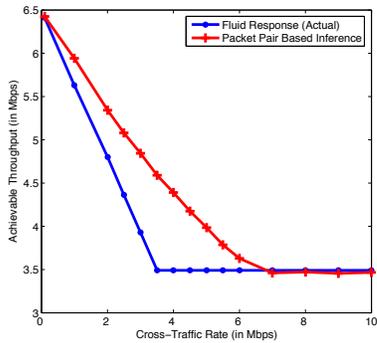


Figure 16: Experimental comparison between packet pair based bandwidth measurements and the actual fluid response in a WLAN link (testbed)

width under FIFO assumptions, follow, instead, the *achievable throughput* when applied to CSMA/CA systems. This idea is illustrated in figure 4 in [25]. There the authors plot the bandwidth estimates gathered using popular tools in an IEEE 802.11 system. The figures show how all the tools used tend to follow the *achievable throughput* rather than the *available bandwidth* when these two metrics start differing in IEEE 802.11 settings.

7.3 Another consequence: packet pair measurements in WLAN links

A common approach to measure the capacity of a network path is the packet-pair technique [23]. Recently, packet pairs have gained momentum as they have been extensively used to develop routing metrics in all-wireless multi-hop networks [21].

However, as a consequence of the results presented in section 6, packet pairs (understood as probes of infinite rate) target the *achievable throughput* when used in a WLAN link. Even more, considering the results presented in section 6, one can see that packet pairs tend to overestimate the value of the *achievable throughput*. Figure 16 illustrates this fact. It plots the actual *achievable throughput* of a WLAN link and the estimation using dispersion measurements of packet-pairs. This is done for different levels of cross-traffic. The *capacity* of the WLAN link is kept constant for all the measurement process at 6.5Mbps (i.e. there are no channel propagation errors). As one can see the packet-pair does not point at the *capacity* in the whole measurement region except when no contending traffic is present.

7.4 An application of results: bandwidth measurement as a simulation warm-up problem

The transient-state of the access delay described in this paper can be understood as a simulation warm-up problem. This is a classical problem in the theory of simulation that has been extensively studied (e.g. [32]).

The literature proposes several techniques to get rid of the effects that samples taken during the transient-state period may induce to the measurement results. A common technique is to enlarge the simulation time in order to assure that transient-state observations are averaged out. This would be

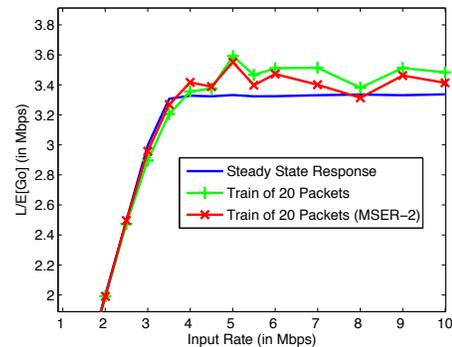


Figure 17: MSER-2 based measurement

equivalent, in our case, to sending longer packet sequences, with the increase of intrusiveness that this entails.

Another technique is trying to infer the duration of the transitory and then truncating the sample sequence. The MSER-m technique is a popular heuristic used to automate the detection of transient-state durations. We have applied this heuristic to our scenario. The idea is to remove from dispersion measurements, those packets that the MSER-m selects as part of the transitory.

Figure 17 illustrates this observation. We apply an MSER-2 heuristic to the inter-arrival time of the packets of a 20 packet train sequence. As the figure shows, when we remove the packets that the heuristic reports as part of the transient-state, the curve better approaches the steady-state behavior. An this is achieved without the need to enlarge the number of packets sent. This could be applied to existing tools [1, 17, 18, 19, 20, 22, 23] in order to improve their accuracy and/or reduce their convergence time.

8. CONCLUSIONS

This paper presents a study of the bandwidth measurement problem when applied to CSMA/CA based systems. On one side the paper presents a complete model of the rate response curve of the system in steady-state completing state-of-art literature related to the topic.

On the other side, the paper analyzes the transient-state behavior of the system considered. This study reveals that the access delay of probing packets undergoes a transitory regime before reaching the steady-state. Additionally, it provides some bounds on the duration of such transient-state regime that can be used, in practice, to design bandwidth measurement tools. Finally, the study provides some insights on how this transient-state regime deviates rate response curves based on short packet trains, and how the effects of this deviation can be safely reduced without increasing the intrusiveness of the measurement process.

9. ACKNOWLEDGEMENTS

This work has been partially funded by the COST Action IC0703 "Data Traffic Monitoring and Analysis", by the Spanish Ministry of Science and Innovation under grant number TEC2008-06826/TEC (project ARTICO), and by the Catalan Regional Government under grants 2009SGR-1140 and 2009SGR-940.

10. REFERENCES

- [1] R.S.Prasad et al. "Bandwidth Estimation: Metrics, Measurement Techniques and Tools", in Proc. of IEEE Network, 2003
- [2] Mingzhe Li et al. "Packet Dispersion in IEEE 802.11 Wireless Networks", in Proc. of IEEE Local Computer Networks, 2006
- [3] Karthik Lakshminarayanan, "Bandwidth Estimation in Broadband Access Networks" In Proc. of the ACM/SIG IMC, 2004
- [4] Mingzhe Li, Mark Claypool, and Robert Kinicki, "Packet Dispersion in IEEE 802.11 Wireless Networks", in Proc. of PAEWN, 2007
- [5] Andreas Johnsson et al. "An Analysis of active end-to-end bandwidth measurements in wireless networks" In Proc. of IEEE E2EMon, 2006
- [6] Issariyakul T. et al., "Exact Distribution of access delay in IEEE 802.11 DCF MAC", In Proc. of IEEE GLOBECOM, 2005
- [7] Shahrokh Valaee, "Bandwidth Estimation and Distributed Traffic Regulation in Wireless Local Area Networks" In Proc. of PIMRC, 2008
- [8] G. Bianchi, "Performance Analysis of the IEEE 802.11 Distributed Coordination Function" in Proc. of IEEE JSAC, 2000
- [9] Marc Portoles-Comeras et al, "EXTREME: Combining the ease of management of multi-user experimental facilities and the flexibility of proof of concept testbeds" In proc. of the IEEE TridentCom, 2006
- [10] Multi-Generator (MGEN) (online)
<http://cs.itd.nrl.navy.mil/work/mgen/>
- [11] Attila Pásztor and Darryl Veitch, "PC Based precision timing without GPS" in Proc. of ACM Sigmetrics, 2002
- [12] The Network Simulator NS2 (online)
<http://www.isi.edu/nsnam/ns/>
- [13] Melander B., Björkman M., Gunningberg P. "Regression based available bandwidth measurements" in Proc. of SPECTS, 2002
- [14] Liu X., Ravindran K., Liu B., and Loguinov, D., "A Queuing-Theoretic Foundation of Available Bandwidth Estimation: Single-Hop Analysis" in IEEE/ACM Transactions on Networking, 2007
- [15] X. Liu, K. Ravindran, and D. Loguinov, "Multi-Hop Probing Asymptotics in Available Bandwidth Estimation: Stochastic Analysis," In Proc. of ACM/SIG IMC, 2005
- [16] NIST/SEMATECH e-Handbook of Statistical Methods (online)
<http://www.itl.nist.gov/div898/handbook/> (Sec. 1.3.5.6)
- [17] M. Jain et al. "End-to-End Available Bandwidth: Measurement Methodology, Dynamics and Relation with TCP Throughput", In Proc. of ACM Sigcomm, 2002
- [18] N. Huet et al. "Evaluation and Characterization of Available Bandwidth Techniques" in IEEE JSAC, 2003
- [19] V. J. Ribeiro, R. H. Riedi, R. G. Baraniuk, J. Navratil, and L. Cottrell, "pathChirp: Efficient available bandwidth Estimation for Network Paths" In Proc. of PAM, 2003
- [20] J. Strauss, "A Measurement Study of available bandwidth Estimation tools" In Proc. of the ACM/SIG IMC, 2003
- [21] R. Draves et al. "Routing in Multi-Radio, Multi-Hop Wireless Mesh Networks" In Proc. of ACM Mobicom, 2004
- [22] Albert Cabellos-Aparicio et al., "A Novel Available Bandwidth Estimation and Tracking Algorithm", in Proc. of IEEE E2EMon, 2008
- [23] Dovrolis C., Ramanathan P., and Moore, D. "Packet dispersion techniques and a capacity estimation methodology" in IEEE/ACM Transaction on Networking, 2004
- [24] Marc Portoles-Comeras et al., "Framework for characterizing hardware deployed in wireless mesh networking testbeds" In Proc. of Tridentcom 2007
- [25] Michael Bredel and Markus Fidler, "A Measurement Study of Bandwidth Estimation in IEEE 802.11g Wireless LANs using the DCF" in Proc. of IFIP Networking 2008
- [26] Ren Wang, Kenshin Yamada, M. Yahya Sanadidi, and Mario Gerla, "TCP with sender-side intelligence to handle dynamic, large, leaky pipes" in IEEE JSAC, 2005
- [27] V. Konda and J. Kaur, "RAPID: Shrinking the Congestion Control Timescale", in Proc. of IEEE INFOCOM, 2009
- [28] Hu, n. et al, "Improving tcp startup performance using active measurements: Algorithm and evaluation", in Proc. of IEEE ICNP, 2003
- [29] Yong Zhu, Constantinos Dovrolis, Mostafa Ammar, "Dynamic overlay routing based on available bandwidth estimation: a simulation study" in Elsevier Computer Networks, 2006
- [30] Robert Carter, Mark Crovella, "Dynamic Server Selection using Bandwidth Probing in Wide-Area Networks", in Boston University, Technical Report, 1996
- [31] Cao Le Thanh Man et al, "ImTCP: TCP with an inline measurement mechanism for available bandwidth" in Elsevier Computer Communications, 2005
- [32] J. A. Joines, Jr., R. R. Barton, et al., "A Comparison Of Five Steady-State Truncation Heuristics For Simulation", in Proc. of Winter Simulation Conference, 2000

APPENDIX

A. VALIDATION SETUP

This section introduces the simulation and experimentation settings used to gather measurement data and validate theoretical findings. Experimentation has been carried out within the EXTREME framework (see [9]). This is a multi-purpose networking experimental platform. The main advantage of this platform is its high automation capabilities that allow automatic execution, data collection and data processing of several repetitions of an experiment.

The WLAN devices used are Z-COM ZDC XI-626 cards which carry the popular Prism chipset. These wireless devices are controlled using computer nodes of the EXTREME cluster. In all cases these nodes are Pentium IV PCs with a 3GHz processor, 512MB of RAM memory and running

Linux OS, with kernel 2.4.26. To control these devices, the EXTREME automation system makes use of the wireless extensions API.

In order to generate the traffic (probing and cross-traffic), we make use of the Multi-GENerator toolset [10]. However, in order to increase the accuracy of the time-stamping procedure, both at sender and receiver sides, network device drivers have been conveniently modified to timestamp packets just before they are laid down to the hardware (sending side) and just after getting them from the hardware (receiving side). This follows some of the ideas described in [11].

Figure 2 shows the basic setup used throughout the section for experimentation. The probing traffic is sent between two stations that are conveniently synchronized. This synchronization is achieved by sending frequent NTP updates through a parallel wired interface between the NTP server and the measurement nodes. Using this method we achieve accuracies of delay measurement in the order of ten microseconds.

Some of the experiments required a large amount of repetitions to achieve accurate convergence of results. Since this is difficult to achieve in a testbed we have also used a simulator. Specifically we have replicated the testbed (figure 2) using NS2 (ver. 2.29 [12]). The main difference between the testbed and the simulator is that the latter includes scenar-

ios with up to 5 contending nodes. Following some recent research results [24] both the testbed and the simulator went through a thorough calibration process in order to assure that the results gathered are comparable.

The simulator uses the NO Ad-Hoc Routing Agent. This agent supports static routing configurations over wireless networks and does not send any routing related packets. This avoids possible interferences with probe or cross-traffic. Regarding the configuration, all the experiments use the default MAC and PHY 802.11 layers included into the NS2 package. The queues used are infinite, this way we avoid dealing with packet losses, which are irrelevant for our study. Finally all the wireless nodes are static and equally spaced from the Access Point. The physical transmission rate is set to 11Mbps and RTS/CTS is not used.

Finally, we have also developed a queuing simulator using Matlab. The motivation for this is that the probing process in a WLAN presents multiple components that are difficult to isolate from each other in an experimentation setting or even through simulations. The queuing simulator convolves a series of packet arrivals with a series of service times in order to measure several metrics such as the queuing length distribution and the output dispersion (inter-arrival) of packets. The input parameters are gathered from experimentation measurements in order to keep the results as close to the real behavior as possible.